OPTIMAL DESIGNS FOR LONGITUDINAL AND FUNCTIONAL DATA[1]

April 2016

Second Revision

Hao Ji[2]

Department of Statistics

University of California, Davis

One Shields Avenue

Davis, CA 95616 U.S.A.

Phone: 1 (530) 400-5942

Fax: 1 (530) 752-7099

Email: haoji@ucdavis.edu

Hans-Georg Müller

Department of Statistics

University of California, Davis

One Shields Avenue

Davis, CA 95616 U.S.A.

Phone: 1 (530) 752-2361

Fax: 1 (530) 752-7099

Email: mueller@wald.ucdavis.edu

arXiv:1604.05375v1 [stat.ME] 18 Apr 2016

[1]Research supported by NSF grants DMS-1104426 and DMS-1407852

[2]Corresponding Author

ABSTRACT

We propose novel optimal designs for longitudinal data for the common situation where the resources for longitudinal data collection are limited, by determining the optimal locations in time where measurements should be taken. As for all optimal designs, some prior information is needed to implement the proposed optimal designs. We demonstrate that this prior information may come from a pilot longitudinal study that has irregularly measured and noisy measurements, where for each subject one has available a small random number of repeated measurements that are randomly located on the domain. A second possibility of interest is that a pilot study consists of densely measured functional data and one intends to take only a few measurements at strategically placed locations in the domain for the future collection of similar data. We construct optimal designs by targeting two criteria:

(a) Optimal designs to recover the unknown underlying smooth random trajectory for each subject from a few optimally placed measurements such that squared prediction errors are minimized;

(b) Optimal designs that minimize prediction errors for functional linear regression with functional or longitudinal predictors and scalar responses, again from a few optimally placed measurements.

The proposed optimal designs address the need for sparse data collection when planning longitudinal studies, by taking advantage of the close connections between longitudinal and functional data analysis. We demonstrate in simulations that the proposed designs perform considerably better than randomly chosen design points and include a motivating data example from the Baltimore longitudinal study of aging. The proposed designs are shown to have an asymptotic optimality property.

KEY WORDS: Asymptotics, Coefficient of Determination, Functional Data Analysis, Functional Principal Components, Gaussian Process, Karhunen-Loève Expansion, Longitudinal Data, Prediction Error, Sparse Design.

1. Introduction

Functional data analysis has become increasingly useful in various fields. In many applications, especially in longitudinal studies, often only a few repeated measurements can be obtained for each subject or item, due to cost or logistical constraints that limit the number of measurements. In some functional/longitudinal data where the recordings are sparse and have been taken at irregular time points, functional data analysis methodology has proved useful to infer covariance structure and trajectories (Staniswalis and Lee 1998; Yao *et al.* 2005a; Li and Hsing 2010). While ideally, longitudinal and functional data would be measured on a dense grid, in practical studies one usually encounters constraints on data collection. It is then of interest to have criteria and principles to determine where on the domain (usually but not necessarily a time interval) one should place a given number of measurements so as to minimize prediction errors when recovering the unobserved trajectories for each subject or to predict a response that is associated with each longitudinal trajectory. Sparse sampling is also of interest in applications where one has available densely sampled functional data, but can sample at only a few important locations in future longitudinal data collection. The question we address here is where these locations should be. Answering these questions is for example of keen interest to determine optimal monitoring schedules for children in developing countries, aiming to recover their growth trajectories from sparse measurements in order to assess growth stunting and faltering.

Several previous studies have discussed methods and algorithms for finding optimal designs in the dense functional data case where data are sampled on a dense regular grid (Ferraty *et al.* 2010; Delaigle *et al.* 2012), or belong to a special family of functions (McKeague and Sen 2010), where the emphasis has been on nonparametric functional regression and classification. However, these methods cannot be extended to the case of sparse functional/longitudinal data, which is a case of crucial interest, as the design selection impacts the planning of longitudinal studies that are often cost-intensive. A crucial design feature is where to place a limited number of future longitudinal observations. While the connections between longitudinal and nonparametric approaches are being increasingly studied (Guo 2004; Xiang *et al.* 2013) and there are also studies on designs for classical parametric longitudinal models (Mentre *et al.* 1997; Anisimov *et al.* 2007) and for random processes and fields (Zagoraiou and Baldi Antognini 2009; Fedorov and Leonov 2013), to our knowledge, there is no previous work on optimal designs for longitudinal studies when the underlying longitudinal trajectories are viewed as smooth random functions. The proposed optimal designs fill this gap and

are found to perform well in simulations (see Section 5).

Our approach is motivated by the idea of design selection for longitudinal studies such as the well-known Baltimore Longitudinal Study of Aging (BLSA) (Shock *et al.* 1984; Pearson *et al.* 1997), where (among other variables) body mass index profiles are measured sparsely in time. Currently available data essentially feature random timing of the measurements (see Figure 1). Our methods will be useful to determine optimal designs for a follow-up study or to recruit new subjects, where we develop optimal designs for (1) recovering the unknown smooth underlying trajectories, as these cannot be easily obtained with nonparametric methods due to the sparseness of the measurements, and (2) for predicting scalar responses in functional linear models. As we will demonstrate, sparsely and irregularly sampled data from a pilot study suffice to construct consistent estimates for the optimal designs. A design resulting from our methodology is depicted in Figure 1, where the three highlighted locations are optimally placed to recover the underlying smooth BMI trajectories when one is constrained to select just three measurements of BMI in future studies.

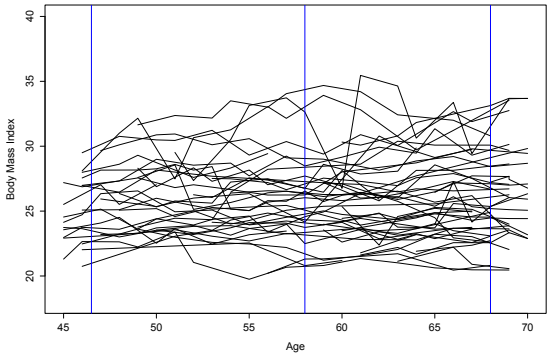

Figure 1: Spaghetti plot for a subset of body mass index (BMI) profiles observed in the Baltimore Longitudinal Study of Aging. The blue vertical bars denote the estimated locations of the optimal design points for recovering BMI trajectories, among designs with three measurement locations.

The vector of function values $\mathbf{X} = (X(t_1), \ldots, X(t_p))^T$ of a generic process X at design points $\mathbf{t} = (t_1, \ldots, t_p)^T$, where ^{T} denotes transpose, is not observed in practice, as observations are contaminated with measurement errors. We refer to the observed values at the design points as $\mathbf{U} = (U(t_1), \ldots, U(t_p))^T$, where

$$U(t_j) = X(t_j) + e_j, \quad Ee_j = 0, \text{ Var}(e_j) = \sigma^2, \; j = 1, \ldots, p, \tag{1}$$

and errors e_j are independent. For trajectory recovery, we use best linear predictors that we denote by B (Rice and Wu 2001)

$$B(X(t)|\mathbf{U}) = \mu(t) + \text{Cov}(X(t), \mathbf{U})\text{Cov}(\mathbf{U})^{-1}(\mathbf{U} - \boldsymbol{\mu}), \quad \mu(t) = E(X(t)), \tag{2}$$

with $\boldsymbol{\mu} = E\mathbf{U} = (\mu(t_1), \mu(t_2), ..., \mu(t_p))^T$; for details of the derivation see Online Supplement A.1. Under Gaussian assumptions these are also the best predictors, as then $E(X(t)|\mathbf{U}) = B(X(t)|\mathbf{U})$. Optimal designs for trajectory recovery are derived by minimizing the expected squared distance between these best linear predictors and the true trajectories, which turns out to be equivalent to maximizing a generalized coefficient of determination with respect to the design points (for details see Online Supplement A.3).

For response prediction, where a functional predictor is coupled with a scalar response, the functional linear model is a classical approach (Cardot *et al.* 1999; Ramsay and Silverman 2005),

$$E(Y|X) = \mu_Y + \int_{\mathcal{T}} \beta(t) X^c(t) dt, \tag{3}$$

where $\mu_Y = EY$, $\beta(t)$ is the regression coefficient function, and $X^c(t)$ is the centered predictor process, i.e. $X^c(t) = X(t) - \mu(t)$. The difficulty in longitudinal designs is that the integral on the r.h.s of (3) cannot be evaluated, whenever there is sparse sampling of trajectories X_i. Therefore, for sparse designs, we condition on $\mathbf{U}$ on both sides of the model in (3), giving the sparse version

$$B\{E(Y|X)|\mathbf{U}\} = \mu_Y + \int_{\mathcal{T}} \beta(t) B(X^c(t)|\mathbf{U}) dt, \tag{4}$$

where again, under the Gaussian assumption, $E(X^c(t)|\mathbf{U}) = B(X^c(t)|\mathbf{U})$ and $E\{E(Y|X)|\mathbf{U}\} = B\{E(Y|X)|\mathbf{U}\}$. Then the optimal designs $\mathbf{t} = \{t_1, \ldots, t_p\}$ are defined as the minimizers of the squared prediction error $E(Y - B\{E(Y|X)|\mathbf{U}\})^2$.

The paper is organized as follows. We discuss optimal designs for trajectory recovery and response prediction in Sections 2 and 3, respectively, followed by estimated optimal designs in section 4 and numerical implementations in Section 5, with simulation studies in Section 6. Data analysis examples from various areas are in Section 7, and asymptotic results in Section 8.

2. Optimal Designs for Trajectory Recovery

For fixed p, consider a generic set of non-random design points $\mathbf{t} = (t_1, \ldots, t_p)^T$, and corresponding values of the underlying process, $\mathbf{X} = (X(t_1), \ldots, X(t_p))^T$, with noisy observations

$\mathbf{U} = (U(t_1), \ldots, U(t_p))^T$ and mean vector $\boldsymbol{\mu} = (\mu(t_1), \ldots, \mu(t_p))^T$. We aim to find optimal designs $\mathbf{t}_{\mathbf{X}}^*$ with respect to a target criterion. For the underlying process X we write $\mu(t) = EX(t)$ for the mean and $\Gamma(s,t) = \text{cov}(X(s), X(t))$ for the auto-covariance function, and the covariance matrices $\boldsymbol{\Gamma} = \text{Cov}(\mathbf{X})$ and $\boldsymbol{\Gamma}_* = \text{Cov}(\mathbf{U})$, where we note that (1) implies $\boldsymbol{\Gamma}_* = \boldsymbol{\Gamma} + \sigma^2 I_p$, with I_p denoting the $p \times p$ identity matrix. With $\mu(t) = EX(t)$, the best linear predictor in (2) becomes

$$\begin{aligned} B(X(t)|\mathbf{U}) &= q(t) + \boldsymbol{\alpha}(t)^T \mathbf{U}, \qquad q(t) = \mu(t) - \boldsymbol{\gamma}(t)^T \boldsymbol{\Gamma}_*^{-1} \boldsymbol{\mu}, \\ \boldsymbol{\alpha}(t) &= \boldsymbol{\Gamma}_*^{-1} \boldsymbol{\gamma}(t), \qquad \boldsymbol{\gamma}(t) = \text{Cov}(\mathbf{U}, X(t)) = (\Gamma(t_1, t), \ldots, \Gamma(t_p, t))^T, \end{aligned} \tag{5}$$

where $\boldsymbol{\gamma}(t)$ is a p-dimensional vector of covariances associated with the non-random time points $\mathbf{t} = (t_1, \ldots, t_p)$. Here and in the following, expectations and covariances are considered to be conditional on designs $\mathbf{t}$.

Our goal is to minimize the mean integrated squared error (MISE) of recovered trajectories that are obtained with the best linear predictors $B(X(t)|\mathbf{U})$ as a function of the design points $\mathbf{t} = (t_1, \ldots, t_p)^T$ that exert their influence through $\mathbf{X} = (X(t_1), \ldots, X(t_p))^T$, i.e. to minimize

$$\text{MISE}(\mathbf{t}) = E \int_{\mathcal{T}} [X(t) - B(X(t)|\mathbf{U})]^2 dt. \tag{6}$$

The performance of recovered trajectories at a fixed point $t \in \mathcal{T}$ can be quantified by a point-wise coefficient of determination, defined as

$$R^2(t) = \frac{\text{Var}(B(X(t))|\mathbf{U})}{\text{Var}(X(t))} \tag{7}$$

This motivates to maximize an overall coefficient of determination with respect to the design points. It is easy to see that

$$R_X^2 = \frac{\int_{\mathcal{T}} \boldsymbol{\gamma}(t)^T \boldsymbol{\Gamma}_*^{-1} \boldsymbol{\gamma}(t) dt}{\int_{\mathcal{T}} \text{Var}(X(t)) dt}, \tag{8}$$

is a well-defined coefficient of determination for all $\mathbf{t} \in \mathcal{T}(\delta_0)$ for some $\delta_0 > 0$, where $\mathcal{T}(\delta_0)$ is defined as

$$\mathcal{T}(\delta_0) = \{(t_1, \ldots, t_p)^T \in \mathcal{T}^p, \ \text{Cov}(U(t_1), \ldots, U(t_p)) - \delta_0 I_p \ \text{ is positive definite}\}. \tag{9}$$

Requiring $\mathbf{t} \in \mathcal{T}(\delta_0)$ ensures uniform matrix invertibility across designs. We assume that the true optimal design also satisfies $\mathbf{t}_X^0 \in \mathcal{T}(\delta_0)$. In Online Supplement A.3 we show that minimizing the MISE in dependence on $\mathbf{t}$ as in (6) is equivalent to maximizing R_X^2 in (8), which is in turn equivalent to find

$$\mathbf{t}_X^* = (t_{X_1}^*, \ldots, t_{X_p}^*)^T = \text{argmax}_{(t_1, \ldots, t_p)^T \in \mathcal{T}(\delta_0)} \int_{\mathcal{T}} \boldsymbol{\gamma}(t)^T \boldsymbol{\Gamma}_*^{-1} \boldsymbol{\gamma}(t) dt. \tag{10}$$

3. Optimal Designs for Predicting Scalar Responses

We assume here the same setting and conditions as in the previous section and aim at finding optimal designs $\mathbf{t}_Y^*$ with respect to a target criterion that is specific for the functional linear model in (3).

The best linear predictor in (4), using (5), is seen to be

$$\begin{aligned} B\{E(Y|X)|\mathbf{U}\} &= \mu_Y + \int_{\mathcal{T}} \beta(t) B(X^c(t)|\mathbf{U}) dt \\ &= \mu_Y + \int_{\mathcal{T}} \beta(t) \boldsymbol{\alpha}(t)^T (\mathbf{U} - \boldsymbol{\mu}) dt = \mu_Y + \boldsymbol{\beta}_p^T (\mathbf{U} - \boldsymbol{\mu}), \end{aligned} \tag{11}$$

with $\boldsymbol{\beta}_p = \int_{\mathcal{T}} \beta(t)\boldsymbol{\alpha}(t)dt$. A basic tool for the following derivations is the Karhunen-Loève expansion of square integrable random processes,

$$X(t) = \mu(t) + \sum_{k=1}^{\infty} \zeta_k \psi_k(t), \tag{12}$$

with eigenfunctions ψ_k, $k = 1, 2, \ldots$ of the covariance operator of X, and uncorrelated (independent in the Gaussian case) functional principal components (FPCs) ζ_k, $k = 1, 2, \ldots$. The eigenfunctions form an orthonormal basis of the space $\mathcal{L}^2(\mathcal{T})$ and one has the covariance expansion

$$\Gamma(s,t) = \text{Cov}(X(s), X(t)) = \sum_{k=1}^{\infty} \rho_k \psi_k(s) \psi_k(t), \tag{13}$$

where the eigenvalues ρ_k of the covariance operator are positive and ordered, $\rho_1 > \rho_2 > \ldots$. The FPCs satisfy $E(\zeta_k) = 0$ and $\text{Var}(\zeta_k) = \rho_k$ for all k.

The regression coefficient function $\beta(t)$ of the FLM (3) can be expanded in the eigenbasis representation (He *et al.* 2000), with convergence under mild regularity conditions,

$$\beta(t) = \sum_{k=1}^{\infty} \frac{E(\zeta_k Y)}{E(\zeta_k^2)} \psi_k(t). \tag{14}$$

Observing that for the cross-covariance function

$$C(t) = \text{Cov}(X(t), Y) = \text{Cov}(\sum_{k=1}^{\infty} \zeta_k \psi_k(t), Y) = \sum_{k=1}^{\infty} E(\zeta_k Y) \psi_k(t) \tag{15}$$

and using the orthonormality of the eigenfunctions ψ_k, it is easy to see that $\sigma_k = E(\zeta_k Y)$ can be written as

$$\sigma_k = \int_{\mathcal{T}} C(t) \psi_k(t) dt. \tag{16}$$

By (5), (14) and (16), we have with $\boldsymbol{\psi}_k = (\psi_k(t_1), ..., \psi_k(t_p))^T$ and $\mathbf{C} = (C(t_1), ..., C(t_p))^T$,

$$\begin{aligned} \boldsymbol{\beta}_p &= \int_{\mathcal{T}} \beta(t)\boldsymbol{\alpha}(t)dt = \boldsymbol{\Gamma}_*^{-1} \int_{\mathcal{T}} (\sum_{k=1}^{\infty} \frac{\sigma_k}{\rho_k}\psi_k(t))(\sum_{k=1}^{\infty} \rho_k \psi_k(t)\boldsymbol{\psi}_k)dt \\ &= \boldsymbol{\Gamma}_*^{-1} \sum_{k=1}^{\infty} \sigma_k \boldsymbol{\psi}_k = \boldsymbol{\Gamma}_*^{-1}\mathbf{C}. \end{aligned} \tag{17}$$

Similar to trajectory recovery in the last section, we propose to minimize the prediction error

$$E[Y - B\{E(Y|X)|\mathbf{U}\}]^2 = E[Y - \mu_Y - \boldsymbol{\beta}_p^T(\mathbf{U} - \boldsymbol{\mu}_{\mathbf{X}})]^2. \tag{18}$$

This is shown in Online Supplement A.3 to be equivalent to maximizing the following coefficient of determination R_Y^2 that quantifies prediction power,

$$R_Y^2 = \frac{\text{Var}[B\{E(Y|X)|\mathbf{U}\}]}{\text{Var}(Y)}, \tag{19}$$

where we assume that the true optimal designs for regression case $\mathbf{t}_Y^0$ lies in $\mathcal{T}(\delta_0)$.

To maximize R_Y^2, it is equivalent to find

$$\mathbf{t}_Y^* = \text{argmax}_{(t_1,...,t_p)^T \in \mathcal{T}(\delta_0)} \text{Var}[B\{E(Y|X)|\mathbf{U}\}] = \text{argmax}_{(t_1,...,t_p)^T \in \mathcal{T}(\delta_0)} \boldsymbol{\beta}_p^T \boldsymbol{\Gamma}_* \boldsymbol{\beta}_p.$$

Therefore, by (17), the optimization criterion can be simplified to

$$\mathbf{t}_Y^* = \text{argmax}_{(t_1,...,t_p)^T \in \mathcal{T}(\delta_0)} \mathbf{C}^T \boldsymbol{\Gamma}_*^{-1} \mathbf{C}. \tag{20}$$

4. Estimated Optimal Designs

While the population optimal designs were derived in the previous sections, in practice they must be estimated from available data. The available observations are

$$U_{ij} = X_i(t_{ij}) + e_{ij},\ 1 \le i \le n,\ 1 \le j \le m_i, \tag{21}$$

$$Y_i = \mu_Y + \int_{\mathcal{T}} \beta(t) X_i^c(t)dt + \epsilon_i,\ 1 \le i \le n, \tag{22}$$

where (22) only applies to the prediction scenario. Here $(X_i, Y_i),\ i = 1, \ldots, n$, are independent realizations of (X, Y), with m_i the number of observed function values for each subject, the t_{ij} are randomly located time points on $\mathcal{T}$ with density function $f_T(\cdot)$, and the e_{ij} and the ϵ_{ij} are random errors with zero mean and variance σ^2 and σ_Y^2, respectively. We assume throughout that the (X_i, Y_i), the e_{ij} and the ϵ_i are all independent. A notable feature of this data model is that it includes noise not only in the responses Y_i but also in the recordings of the random trajectories.

From the data, estimates of mean function $\mu(t)$, auto-covariance function $\Gamma(s,t)$ and cross-covariance function $C(t)$ are obtained on a user-defined fine grid covering $\mathcal{T}$ and these are denoted as $\widehat{\mu}(t)$, $\widehat{\Gamma}(s,t)$ and $\widehat{C}(t)$, respectively. Consistent estimates of these quantities from the pilot study are needed to obtain consistent estimates of the optimal designs. For densely observed functional data, cross-sectional estimates are suffcient. Methods to overcome the difficulty of sparse sampling when targeting the mean and covariance functions have been addressed by various authors (Yao *et al.* 2005a,b; Staniswalis and Lee 1998; Li and Hsing 2010). The sparsely sampled case is different from the more commonly considered situation of densely sampled functional data, where individual curves can be consistently estimated by direct smoothing (Rice 2004) and the covariance function is readily estimated by cross-sectional averaging (Ramsay and Silverman 2005).

In the sparse case, these direct approaches do not lead to consistent estimates, due to the sparseness and lack of balance of the measurements. The way forward is to pool data for estimating mean and covariance functions, borrowing strength from the entire sample. For the required smoothing steps, we adopt one- and two-dimensional local linear smoothing, with further details in the following section. The estimated auto-covariance function is further regularized by retaining only the positive eigenvalues and eigenvectors of the smoothed covariance function, so that $\widehat{\Gamma}(s,t)$ is non-negative definite. An estimate $\widehat{\sigma}^2$ of σ^2 is also needed and this is discussed in Section 5. For any p-dimensional vector $(t_1,...,t_p)^T$ of design points picked from the user-specified dense grid, we then have estimates $\widehat{\boldsymbol{\mu}}_X$, $\widehat{\mathbf{\Gamma}}_*$, $\widehat{\boldsymbol{\gamma}}$ and $\widehat{\mathbf{C}}$ for $\boldsymbol{\mu}$, $\mathbf{\Gamma}_*$, $\boldsymbol{\gamma}$ and $\mathbf{C}$, respectively. Then the estimated optimal designs are:

1. For trajectory recovery:

$$\widehat{\mathbf{t}}_X^* = \text{argmax}_{(t_1,...,t_p)^T} \int_{\mathcal{T}} \widehat{\boldsymbol{\gamma}}^T(t) \widehat{\mathbf{\Gamma}}_*^{-1} \widehat{\boldsymbol{\gamma}}(t) dt, \tag{23}$$

   where all integrals are implemented with trapezoidal integration.

2. For scalar response regression:

$$\widehat{\mathbf{t}}_Y^* = \text{argmax}_{(t_1,...,t_p)^T} \widehat{\mathbf{C}}^T \widehat{\mathbf{\Gamma}}_*^{-1} \widehat{\mathbf{C}}. \tag{24}$$

5. Numerical Implementation

5.1 Mean and Covariance Estimation via Smoothing

First pooling sparse longitudinal data across subjects, we apply local linear estimators (Li and Hsing 2010) to the resulting scatterplots, which depend on a bandwidth h as a tuning (smooth-

ing) parameter. Writing $S_p(t,(Q_j,V_j)_{j=1,\ldots,m},h)$ for a local linear q-dimensional smoother (with $q=1$ or $q=2$) with output at the predictor level t and employing bandwidth h to smooth the scatterplot (Q_j,V_j), where $Q_j \in \mathcal{R}^q$, we obtain estimates $\hat{\mu}(t)$ for the mean function $\mu(t)$ as $S_1(t,(t_{ij},U_{ij})_{i=1,\ldots,n,j=1,\ldots,m_i},h_\mu)$, $t \in \mathcal{T}$, smoothing estimates $\tilde{\Gamma}(s,t)$ for the autocovariance function $\Gamma(s,t)$ as $S_2((s,t),(t_{ij},t_{ik},U_{ij}U_{ik})_{i=1,\ldots,n,j,k=1,\ldots,m_i,j\neq k},h_R)-\widehat{\mu}(s)\widehat{\mu}(t)$, $s,t\in\mathcal{T}$, and estimates $\hat{C}(t)$ for the cross-covariance function $C(t)$ as $S_1(t,(t_{ij},Y_iU_{ij})_{i=1,\ldots,n,j=1,\ldots,m_i},h_S)-\widehat{\mu}(t)\widehat{\mu}_Y$, $t\in\mathcal{T}$. We also obtain the estimate $\widehat{\mu}_Y$ as the sample mean of the scalar responses.

Bandwidths for all smoothing steps are selected by cross validation or generalized cross validation. Details on the smoothing steps are in Online Supplement A.2, and assumptions for establishing consistency of the above smoothing steps in the longitudinal data context are provided in Online Supplement A.4. We used the function FPCA (Yao *et al.* 2005a) in PACE, a free Matlab package (http://www.stat.ucdavis.edu/PACE/) for smoothing and estimating the model components.

From the estimates $\tilde{\Gamma}(s,t)$ we then obtain estimates $\hat{\rho}_j$ and $\hat{\psi}_j$ for eigenvalues and eigenfunctions of predictor processes X by discretization and matrix spectral decomposition. The final autocovariance estimates $\widehat{\Gamma}(s,t)$ are obtained by projecting on the space of non-negative and symmetric surfaces, simply by retaining only the positive eigenvalues and their corresponding eigenvectors (Hall *et al.* 2008), yielding $\widehat{\Gamma}(s,t)=\sum_{j=1,\widehat{\rho}_j>0}^{K}\widehat{\rho}_j\widehat{\psi}_j(s)\widehat{\psi}_j(t)$.

5.2 Stable Covariance Matrix Inversion

As a practical implementation of the matrix inversion condition (9), we apply ridge regression (Hoerl and Kennard 1970), enhancing the diagonal of the autocovariance surface $\widehat{\Gamma}_*(s,t)$, as the matrices that need to be inverted are submatrices of this surface. Adding a suitable ridge parameter $\widehat{\sigma}^2_{new}$ at the diagonal ensures positive definiteness of all relevant p by p submatrices,

$$\widehat{\Gamma}_*(s,t)=\widehat{\Gamma}(s,t)+\sigma^2_{\text{new}}\delta_{st}. \tag{25}$$

Here $\delta_{st}=1$ if and only if $s=t$. The optimization procedures in (23) and (24) are then implemented with $\widehat{\Gamma}_*(s,t)$ or $\widehat{\mathbf{\Gamma}}_*=\widehat{\mathbf{\Gamma}}+\sigma^2_{\text{new}}I_p$. We explored two options to select the ridge parameter σ^2_{new}:

1. Cross-validation. For *trajectory recovery*, target criteria are average root mean squared error, ARE, and relative average root mean squared error, ARE*, defined as

$$ARE \;=\; \frac{1}{n}\sum_{i=1}^{n}\left\{\frac{1}{m_i}\sum_{j=1}^{m_i}(U_i(t_{ij})-\widehat{B}_{-i}(X_i(t_{ij})|\mathbf{U}))^2\right\}^{1/2},$$

$$ARE^* = \sum_{i=1}^{n}\left\{\frac{1}{m_i}\sum_{j=1}^{m_i}(U_i(t_{ij}) - \widehat{B}_{-i}(X_i(t_{ij})|\mathbf{U}))^2\right\}^{1/2} / \sum_{i=1}^{n}\left\{\frac{1}{m_i}\sum_{j=1}^{m_i}U_i(t_{ij})^2\right\}^{1/2} \quad (26)$$

where $(U_i(t_{i1}),\ldots,U_i(t_{im_i}))^T$ and $(\widehat{B}_{-i}(X_i(t_{i1})|\mathbf{U}),\ldots,\widehat{B}_{-i}(X_i(t_{im_i})|\mathbf{U}))^T$ are observed measurements and plugged-in estimated best linear predictors for recovered processes with design $\mathbf{t}$ at the same time points. Leave-one-out versions of ARE and ARE* are easily obtained if one has densely measured functional data in the pilot study, but are usually not viable when the pilot study consists of longitudinal data with sparse measurements, where observed subjects generally will not have recorded measurements at the selected optimal design points.

For *scalar response prediction*, average prediction error (APE) and relative APE are natural criteria that can be easily cross-validated, irrespective of the design of the pilot study,

$$\begin{aligned} APE &= \left\{\frac{1}{n}\sum_{i=1}^{n}(Y_i - \widehat{B}_{-i}\{E(Y_i|X_i)|\mathbf{U}_i\})^2\right\}^{1/2}, \\ APE^* &= \left\{\sum_{i=1}^{n}(Y_i - \widehat{B}_{-i}\{E(Y_i|X_i)|\mathbf{U}_i\})^2\right\}^{1/2} / \left\{\sum_{i=1}^{n}Y_i^2\right\}^{1/2} \end{aligned} \quad (27)$$

where $\widehat{B}_{-i}\{E(Y_i|X_i)|\mathbf{U}_i\}$ is the estimated i-th response obtained with estimated optimal designs that are obtained from a training sample leaving out the i-th observation.

2. Modified Cross-Validation. Direct cross-validation approach is not feasible for the case of sparsely sampled pilot studies. In a modified approach, for each ridge parameter in a candidate set Ω, we repeatedly and randomly partition the training sample $\mathcal{S}$ from the sparsely sampled pilot study into two sets, $\mathcal{S}_A$ and $\mathcal{S}_B$, estimate model components from $\mathcal{S}_A$, and then find the relatively best design $\mathbf{t}^*_{X,\mathcal{S}_B}$ or $\mathbf{t}^*_{Y,\mathcal{S}_B}$ by maximizing the criteria in (10) or (20) over $\mathcal{T}_B$, which is the set of all available designs determined by the random configurations of the design points as observed for the subjects in the sample $\mathcal{S}_B$. We then recover the trajectory or estimate the response for those subjects in $\mathcal{S}_B$ where there is a match of the selected design $\mathbf{t}^*_{X,\mathcal{S}_B}$ or $\mathbf{t}^*_{Y,\mathcal{S}_B}$ with the design for that subject.

Combining L different random partitions, mean ARE or APE are used to evaluate the performance of the ridge parameter choice in (25), yielding the selected parameter

$$\sigma^2_{\text{new}} = \text{argmin}_{\Omega}\sum_{l=1}^{L}\left(\frac{1}{n_B}\sum_{i_B=1}^{n_B}\left\{\frac{1}{m_{i_B}}\sum_{j=1}^{m_{i_B}}(U_{i_B}(t_{i_B,j}) - \widehat{B}_{\mathcal{S}_A}(X_{i_B}(t_{i_B,j})|\mathbf{U}_{i_B}))^2\right\}^{1/2}\right),$$

for trajectory prediction, and

$$\sigma^2_{\text{new}} = \text{argmin}_{\Omega} \sum_{l=1}^{L} \left(\frac{1}{n_B} \sum_{i_B=1}^{n_B} (Y_{i_B} - \widehat{B}_{\mathcal{S}_A}\{E(Y_{i_B}|X_{i_B})|\mathbf{U}_{i_B}\})^2 \right)^{1/2}.$$

for response prediction. Here, i_B is the index of the subjects in $\mathcal{S}_B$ with available measurements at the selected optimal designs $\mathbf{t}^*_{X,\mathcal{S}_B}$ and $\mathbf{t}^*_{Y,\mathcal{S}_B}$ for trajectory recovery and prediction, respectively, with n_B denoting the number of such subjects, and m_{i_B} is the number of measurements for the subject with index i_B. The estimator $\widehat{B}_{\mathcal{S}_A}$ is fitted based on data from sample $\mathcal{S}_A$ only, and index sets $\mathcal{S}_A$ and $\mathcal{S}_B$ depend on the random partition l. This method worked well in simulations and applications with longitudinal pilot studies.

5.3 Sequential Selection of Design Points

Computationally, once the number of design points p has been specified, both trajectory recovery and scalar response prediction involve p-dimensional optimization. Exhaustive search over all combination of grid points is very time consuming when employing optimization algorithms such as simulated annealing (Kirkpatrick *et al.* 1983). A faster alternative is sequential selection, which is a greedy algorithm, where one searches for a global optimal designs when $p = 2$ as an initial step and then adds design points one-by-one iteratively, until the number of design points reaches the desired number. At each step, the target design point is the one that maximizes the selection criteria when adding it to the currently selected design points, which are carried forward unaltered. The sequential method is fast but does not guarantee finding the optimal solution. The performance differences of sequential and exhaustive search selection were found to be relatively small in simulation studies.

6. Simulation Studies

We study the performance of optimal designs for trajectory recovery and scalar response prediction under two separate scenarios. In scenario 1 we consider the case where the pilot study generates densely observed functional data, and in scenario 2 the case where it generates sparse longitudinal data. Random trajectories are generated as $X_i(t) = \mu(t) + \sum_{k=1}^{K} \zeta_{ik}\psi_k(t)$, and observed data as $U_i(t_{ij}) = X_i(t_{ij}) + e_{ij}$, where $e_{ij} \sim N(0, 0.25)$.

For the regression case, the response is chosen as $Y_i = \int \beta(t)X_i(t)dt + \epsilon_i = \zeta_{i,1} - 2\zeta_{i,2} + \zeta_{i,3} - 2\zeta_{i,4} + \epsilon_i$, i.e., as a linear combination of functional principal components of the process, where $\epsilon_i \sim$

$N(0, 0.25)$ are i.i.d. We specify $\mathcal{T} = [0, 10]$, mean function $\mu(t) = \frac{1}{2}t^2 + 2\sin t + 3\cos 2t$, $t \in \mathcal{T}$, and include 10 functional principal components, with eigenvalues $(30, 20, 12, 8, \frac{30}{25}, \frac{30}{36}, \frac{30}{49}, \frac{30}{64}, \frac{30}{81}, \frac{30}{100})^T$, and corresponding eigenfunctions $\psi_k(t) = \sqrt{\frac{2}{10}}\cos\{(\frac{k}{10})\pi t\}$, for $k = 1, 2, \ldots, 10$ and generate data for 100 subjects in the training sample and 1000 in the testing sample. The measurement locations t_{ij} are assumed to form a dense grid for simulation scenario 1, and are sparse, with a random number of 4 to 8 measurement locations per subject for scenario 2. Figure 9 in Online Supplement A.9 shows the Spaghetti plot of the data for the subjects in one training sample.

For both trajectory recovery and prediction for functional linear regression, we applied the proposed procedures for the subjects in the training sample to construct the optimal designs for $p = 2, 3, 4$ with exhaustive search and $p = 2, 3, ..., 8$ with sequential search. Each simulation scenario was repeated 100 times. The optimal ridge parameter σ^2_{new} in (25) was determined by cross-validation for scenario 1 and modified cross-validation for scenario 2 (see subsection 5.2). We compare the median performance with regard to (relative) Average Root Squared Error (ARE) and (relative) Average Prediction Error (APE) defined in (26) and (27) for optimal designs and random designs. These random designs use the same number of design points as the estimated optimal designs, however the locations are sampled from a uniform distribution over all possible locations (we provide further comments on the rationale of comparing with random designs in Online Supplement A.5). The results are summarized in Table 1, and are illustrated in Figure 2 (and also Figure 10 in Online Supplement A.9) for the case where the pilot study is longitudinal with sparsely sampled functional data.

These simulations demonstrate that the proposed optimal designs exhibit better performance than random designs for both trajectory recovery and scalar response prediction, especially for sparse functional data. For trajectory recovery, Figure 10 (in Online Supplement A.9) for sparse pilot designs indicates that recovered trajectories obtained from optimal designs are closer to the underlying true curves than those obtained from median-performance random designs. For scalar response prediction, Figure 2 corroborates the results from Table 1, namely that optimal designs outperform median-performance random designs for the prediction of a subject's response from a few observed measurements only.

The boxplots of ARE and APE for 100 simulation runs in Figure 3 and Figure 10 (in Online Supplement A.9) visualize the variation of performance over the simulations, comparing optimal and median performance random designs, and also show the improvement in performance as the number of design points p increases. For densely observed functional data, there are various penalization

schemes possible to select p, minimizing the sum of ARE (APE) and a penalty that increases with increasing p. However, such schemes are not directly applicable for longitudinal data, because subjects rarely will have been observed at the selected design points for trajectory recovery or prediction. In practice, an upper bound for p frequently will be dictated by cost. Simulation results for the effect of ridge parameter selection on the performance of optimal designs showed that the proposed selection works well (see Online Supplement A.7).

Table 1: Comparing Optimal with Random Designs in 100 Simulations, in terms of mean ARE and relative ARE* (26) (in brackets) for Trajectory Recovery and APE and relative APE* (27) (in brackets) for Response Prediction in a Functional Linear Model. Exhaustive Search used.

Number of Design Points		Mean ARE (ARE*)		Mean APE (APE*)	
	Design	Dense	Sparse	Dense	Sparse
$p = 2$	Optimal	$1.74(.091)$	$1.84(.102)$	$3.37(.272)$	$9.41(.759)$
	Random	$1.91(.101)$	$2.12(.117)$	$10.65(.870)$	$10.79(.906)$
$p = 3$	Optimal	$1.37(.072)$	$1.59(.088)$	$2.79(.224)$	$7.63(.613)$
	Random	$1.65(.087)$	$1.88(.104)$	$8.56(.687)$	$9.80(.793)$
$p = 4$	Optimal	$1.02(.054)$	$1.34(.074)$	$2.54(.204)$	$6.87(.553)$
	Random	$1.46(.077)$	$1.74(.096)$	$6.44(.499)$	$7.65(.616)$

To summarize the simulation results, optimal designs performed very well and in any case better than random designs for both trajectory recovery and scalar response prediction. The costs incurred when adopting the much faster sequential search algorithm care quite small. Unsurprisingly, performance of the optimal designs was seen to improve with increasing number of design points.

7. Data Illustrations

7.1 Mediterranean Fruit Fly Egg-Laying

The Mediterranean fruit fly data, described in (Carey *et al.* 2002), consist of egg-laying profiles for 1000 female Mediterranean fruit flies. For each fly, daily measurements on the number of eggs laid during the day are available from birth to death. A biologically relevant regression problem is to utilize the partial egg-laying profile from day 1 to day 30 to predict the number of eggs that will

be laid during the remaining lifetime for each fly. This yields information about the reproductive potential of the fly at age 30 days, which is related to its evolutionary fitness (Kouloussis *et al.* 2011).

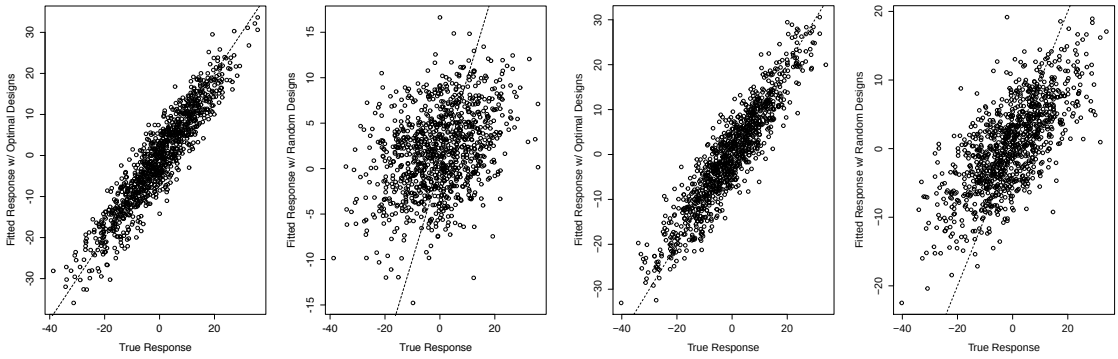

Figure 2: Simulation results for sparsely sampled pilot data. Left two (right two) scatterplots are for fitted versus true responses for predicting a scalar response based on 3 (4) design points. First and third panels (from the left) demonstrate the results for optimal designs, while second and fourth panels (from the left) depict the results for random designs with median performance.

We aim to find the optimal design points for this scalar response regression problem. To prevent censoring, we only include flies that live beyond 30 days. The measurements in the pilot data are dense and regular. Since daily egg-laying counts require constant monitoring of the flies, reducing this task to monitoring of the flies at a few time points is useful to scale up such studies. It is of additional interest to identify key days that are relevant for the prediction of the egg-laying potential. We use the complete available egg-laying profiles to find optimal design points that are most relevant for the prediction of the remaining total number of eggs via a functional linear regression model.

From the 667 subjects surviving more than 30 days, we select a training sample of 500 flies, and a testing sample that consists of the remaining 167 flies. Figure 4 shows the Spaghetti plot for a subset of the training sample. We apply the proposed methodology to find optimal designs for $p = 3$. The relationship between observed and predicted responses is shown in Figure 5. The relative APE (27) is 0.483 when using optimal designs, as opposed to 0.665 using random designs with median performance for $p = 3$. Figure 5 provides a graphical illustration that optimal designs clearly outperform randomly chosen designs with median performance. The three selected optimal design points for $p = 3$ are at days 10, 25 and 26. Their locations are shown as blue vertical lines in Figure 4. These design point locations are consistent with previous findings that both the intensity of

egg-laying at earlier ages and the rate of decline at older ages are closely related to the reproductive potential for individual flies (Müller *et al.* 2001). Two design points are selected on consecutive days 25 and 26, indicating that these locations are useful to gauge the rate of decline in egg-laying, which corresponds to quantifying a derivative in this age range.

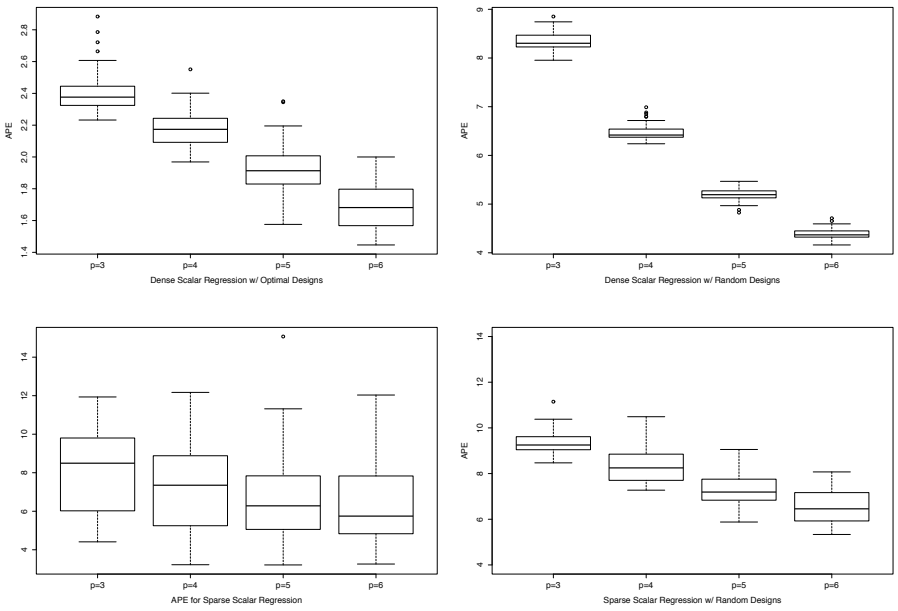

Figure 3: Boxplots of APE from 100 simulation runs for scalar response regression for both dense and sparse scenarios with sequential search algorithm, with $p = 3, 4, 5, 6$. The first (second) row corresponds to the dense (sparse) scenario; the left (right) column has the results for the optimal design (median performance random design).

7.2 The Baltimore Longitudinal Study of Aging

For the Baltimore Longitudinal Study of Aging (BLSA) we aim at identifying optimal designs for recovering Body Mass Index (BMI) profiles from sparse measurements and for predicting a subject's systolic blood pressure (SBP) at old age. To construct optimal designs, we use available pilot data that come from the longitudinal BLSA study, where measurements of BMI are sparse and irregularly spaced. To avoid bias due to censoring effects, we only include subjects with non-missing date of death and for whom age at death is above 70, and consider only the available measurements that were taken within the age range from 45 to 70. We also exclude subjects who had less than four

measurements. The response is taken to be the last SBP measurement before death. Within the study sample, 496 subjects met these criteria and were included in the analysis. A subset of the data, where measurements are connected by straight lines, is shown in Figure 1.

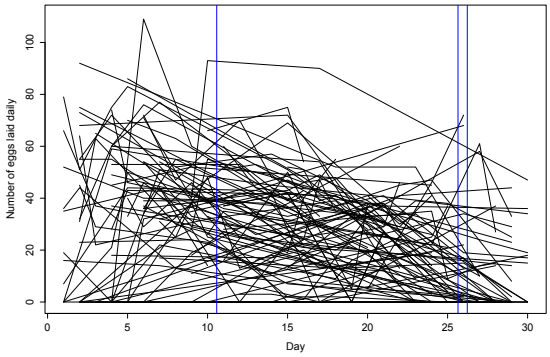

Figure 4: Spaghetti plot for a subset of the training sample for predicting remaining egg-laying potential for female medflies from their egg-laying profiles. The vertical bars denote the locations of the optimal design points for $p = 3$.

A decisive difference between these data and the previous data illustration is that the BMI trajectories that are assumed to generate the observed data are not available, as these pilot data are from a typical longitudinal study with inherently sparse and irregular measurement times. Therefore, only indirect information is available for each subject about the underlying trajectory and only conditional inference about the trajectories is possible, conditioning on the available measurements for each subject (Yao *et al.* 2005a). This is the typical situation one faces when constructing optimal designs in the common situation where the available pilot data are from a longitudinal study. Nevertheless, the construction of optimal designs is still possible since they only depend on the covariance structure of the data, which can still be consistently estimated. To implement the proposed optimal designs for this situation, we use the modified cross validation method to select the ridge parameter.

Here the construction of optimal designs is intended for future data collection in longitudinal studies that will adopt the same fixed optimal design for the subjects to be included in the subsequent study. Due to the fact that the pilot data are coming from a longitudinal study with random measurement locations, subjects included in the pilot study will normally not have measurements at the optimal design locations. As in addition their actual functional trajectories are unknown, it is not

possible within the framework of the BLSA study to directly evaluate the performance of the optimal designs in terms of recovering the BMI trajectories or predicting old age systolic blood pressure (SBP). Alternatively, we consider performance as measured by the coefficients of determination in (8) and (19). The selected optimal designs, along with their associated coefficients of determination, are listed in Table 3 for $p = 3, 4, 5$, for both prediction and trajectory recovery, where exhaustive search was used to determine these designs.

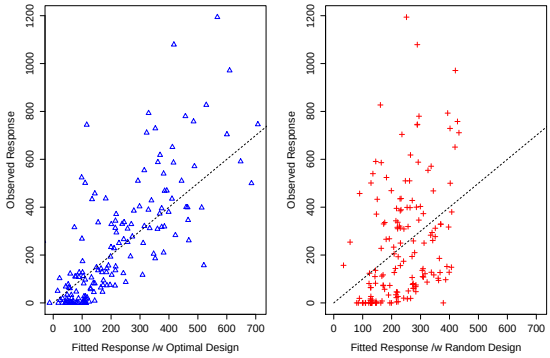

Figure 5: Observed versus fitted (predicted) numbers of eggs that female medflies lay in their remaining lifetime when using $p = 3$ design points in a functional linear regression model. The predictions for optimal designs are in the left panel and those for median performance random designs in the right panel. The black dotted line is the 45 degree line in both graphs.

The following findings are of interest for this and similar situations where one has longitudinal pilot data: First, we the designs for the prediction and the recovery task were found to differ somewhat, especially in terms of older age measurements that are part of the optimal designs for trajectory recovery if $3 \leq p \leq 4$, but are somewhat less relevant for predicting the systolic blood pressure response at 70. Therefore, optimal designs for old age SBP prediction from BMI will not require measurements at older ages. Second, the constructed optimal designs for the same target quantity are consistent in the sense that as p increases, additional design points are added while the previously selected design points are still viable so that selected optimal designs to some extent form a nested sequence, even when exhaustive search is used to determine these designs. Third, the design points for trajectory recovery tend to be more evenly distributed than those for response prediction. In other words, design points that are more or less uniformly distributed across the age range seem to achieve the goal of trajectory recovery well.

The latter point was also seen in the bicycle sharing data, which are described in Online Supplement A.6. The reason behind this might be that the random variation in trends or curvature across subjects is relatively uniform in these examples (see Figure 1). The optimal design points for prediction of old age SBP are generally clustering around earlier ages, which is potentially of interest for public health and prevention. Finally, the coefficients of determination for both trajectory recovery and prediction are increasing with p, showing better performance with increasing number of design points, as was also seen in the simulations. To select the optimal designs together with the optimal p, one can use a suitable penalty for larger p that might reflect the data collection cost in specific applications.

Table 2: Optimal Designs (Ages Where Measurements are Placed) and Associated Coefficients of Determination for Different Numbers of Design Points for the BLSA Data, for Recovering Body Mass Index (BMI) and for Predicting Systolic Blood Pressure (SBP).

	Designs for BMI Trajectory Recovery (R^2_X)	Designs for SBP Prediction (R^2_Y)
$p = 3$	46.5, 58, 68 (0.951)	45, 45.5, 53 (0.863)
$p = 4$	46, 56.5, 62.5, 70 (0.973)	45, 45.5, 52, 52.5 (0.874)
$p = 5$	47, 47.5, 60.5, 61.5, 70 (0.989)	45, 45.5, 53, 53.5, 68.5 (0.886)

Specifically for the BLSA study, comparing designs with different p, a general recommendation would be to take the first measurement at around age 46 or 47, then the second within the age range 58 to 62, and finally a third at around age 70, if $p = 3$ and BMI trajectory recovery is the objective. For predicting the systolic blood pressure, optimal designs are more concentrated in the first half of the age range, and it would be sufficient to take measurements at age 45 , 45.5 and around age 53.

8. Theory

We establish asymptotic consistency and rates of convergence for the proposed optimal designs. Assumptions (A1)-(A7) and the proof of the following main result are in Online Supplement A.4. The rates of convergence that we report here are for the case that the pilot study is longitudinal and generates sparse functional/longitudinal data. In the following, h_μ, h_S and h_R are bandwidths for local linear smoothers for the mean function, cross-covariance function and auto-covariance function, as

specified in (32), (33) and (34) in Online Supplement A.2.

Theorem. *Assume that (A1)-(A7) hold, and that criteria R_X^2 and R_Y^2 are locally concave around the optimal design $\mathbf{t}_X^*$ for trajectory recovery (10) and $\mathbf{t}_Y^*$ for scalar response regression (20) respectively, where $\mathbf{t}_X^*$, $\mathbf{t}_Y^* \in \mathcal{T}(\delta_0)$, defined at (9). For $\theta_n = h_R^2 + \{\log n/(nh_R^2)\}^{1/2}$, and given fixed p, the estimated optimal designs for trajectory recovery $\widehat{\mathbf{t}}_X^*$ given by (10) and for scalar response regression $\widehat{\mathbf{t}}_Y^*$ given by (20) satisfy*

$$||\widehat{\mathbf{t}}_X^* - \mathbf{t}_X^*||_p = O(\theta_n^{1/2}) \ \ a.s. \ \ as \ \ n \to \infty$$

and

$$||\widehat{\mathbf{t}}_Y^* - \mathbf{t}_Y^*||_p = O(\theta_n^{1/2}) \ \ a.s. \ \ as \ \ n \to \infty$$

where $||\mathbf{t} - \mathbf{s}||_p = \max_{1 \le j \le p} |\mathbf{t}_{(j)} - \mathbf{s}_{(j)}|$, where $\mathbf{t}_{(j)}$, $\mathbf{s}_{(j)}$ are the j-th order statistics of designs $\mathbf{t}$ and $\mathbf{s}$. Here, we assume in addition that the smoothing bandwidths satisfy $h_R^2 \lesssim h_\mu \lesssim h_R$, and $h_\mu \sim h_S$, where $a_n \lesssim b_n$ means $a_n = O(b_n)$, and $a_n \sim b_n$ means that $c \le \frac{a_n}{b_n} \le C$ for constants $0 < c < C < \infty$.

This result demonstrates the consistency of the estimated optimal designs, including rates of convergence to the true optimal designs. This provides theoretical justification for our methods. The rate of convergence θ_n is determined by the rate at which the bandwidth h_R for smoothing the cross-covariance surface of the underlying smooth stochastic process converges to 0.

The theorem is proved by first establishing uniform convergence of the optimization criteria (10), (20) with plugged-in estimators of the auxiliary quantities. The second step then is to prove the convergence of the estimated maximizer to the true maximizer. If the pilot study has dense and regular designs, one can apply cross-sectional covariance and mean estimation and by similar arguments provided here obtains analogous results as in the Theorem with the faster rate $\theta_n = n^{-1/2}$.

9. Discussion and Concluding Remarks

We have developed a new method to obtain optimal designs for longitudinal data, by considering such data to be instances of functional data, i.e., by assuming that they are generated from an underlying (but unobservable) smooth stochastic process. This perspective makes it possible to take a pilot sample of sparsely observed longitudinal data and to construct optimal design points for future observations, where optimality refers to various criteria and loss functions.

For the trajectory recovery task, the optimal designs are related to the shapes of the first few eigenfunctions that explain most of the variation. Specifically, the optimal design points are likely to cluster around areas where the variation in the underlying process across subjects is large, and that might correspond to areas where the first few eigenfunctions have peaks or valleys (Hall and Vial 2006), or generally areas where they are large in absolute value. On the other hand, for the scalar regression task, the optimal designs cannot be easily understood from the shape of the eigenfunctions, because these designs depend not only by the auto-covariance, but equally on the cross-covariance between responses and predictor trajectories.

Simulations and data analyses show that optimal designs can lead to substantial gains over random designs for both trajectory recovery and prediction. The constructed optimal designs thus provide guidance for the planning and collection of longitudinal data. The proposed method is conceptually straightforward and the estimating procedure is easy to implement. The method can be applied to a wide range of studies where dense measurements of underlying trajectories would be desirable but cannot be obtained due to various constraints. In various engineering and medical applications, dense measurement designs are too expensive to obtain and therefore sparse designs need to be used. Some further extensions are discussed in online supplement A.8.

While our focus in this paper is on optimal trajectory recovery and optimal prediction as two pertinent and important criteria, there are many other conceivable targets, such as predicting a functional response, or specific features of the underlying random trajectories, for example integrals or derivatives. It is also possible to select a mixed target criterion, targeting both trajectory recovery and optimal prediction. One can then apply the same methodology as we describe here to find the optimal designs for such a mixed target criterion that blends various objectives, with relative weights given to each. For any given design, we can use the proposed criteria to evaluate its relative suitability for both prediction of a response as well as obtaining trajectory estimates in comparison with competing designs, for example by comparing the relative values of ARE (26) and APE (27). This enables evaluation of the relative merits of potentially suboptimal designs that sometimes may be more convenient than optimal designs, or may be the result of extraneous constraints in data collection.

Our methods are also applicable in situations where in principle one can sample functional data densely but in future data collection only plans to sample at a few key locations. We find that the construction of optimal designs for longitudinal studies benefits in many ways from the adoption of a functional approach.

References

Anisimov, V. V., Fedorov, V. V. and Leonov, S. L. (2007) Optimal design of pharmacokinetic studies described by stochastic differential equations. In *mODa 8-Advances in Model-Oriented Design and Analysis*, 9–16. Springer.

Cardot, H., Ferraty, F. and Sarda, P. (1999) Functional linear model. *Statistics and Probability Letters*, **45**, 11–22.

Carey, J. R., Liedo, P., Harshman, L., Zhang, Y., Müller, H.-G., Partridge, L. and Wang, J.-L. (2002) Life history response of mediterranean fruit flies to dietary restriction. *Aging Cell*, **1**, 140–148.

Delaigle, A., Hall, P. and Bathia, N. (2012) Componentwise classification and clustering of functional data. *Biometrika*, **99**, 299–313.

Fanaee-T, H. and Gama, J. (2013) Event labeling combining ensemble detectors and background knowledge. *Progress in Artificial Intelligence*, 1–15.

Fedorov, V. V. and Leonov, S. L. (2013) *Optimal design for nonlinear response models*. CRC Press.

Ferraty, F., Hall, P. and Vieu, P. (2010) Most-predictive design points for functional data predictors. *Biometrika*, **97**, 807–824.

Guo, W. (2004) Functional data analysis in longitudinal settings using smoothing splines. *Statistical Methods in Medical Research*, **13**, 49–62.

Hall, P., Müller, H.-G. and Yao, F. (2008) Modelling sparse generalized longitudinal observations with latent gaussian processes. *Journal of the Royal Statistical Society: Series B*, **70**, 703–723.

Hall, P. and Vial, C. (2006) Assessing extrema of empirical principal component functions. *The Annals of Statistics*, **34**, 1518–1544.

He, G., Müller, H. and Wang, J. (2000) Extending correlation and regression from multivariate to functional data. *Asymptotics in Statistics and Probability*, 197–210.

Hoerl, A. E. and Kennard, R. W. (1970) Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, **12**, 55–67.

Kirkpatrick, S., Gelatt, C. D. and Vecchi, M. P. (1983) Optimization by simulated annealing. *Science*, **220**, 671–680.

Kouloussis, N. A., Papadopoulos, N. T., Katsoyannos, B. I., Müller, H.-G., Wang, J.-L., Su, Y.-R., Molleman, F. and Carey, J. R. (2011) Seasonal trends in ceratitis capitata reproductive potential derived from live-caught females in greece. *Entomologia Experimentalis et Applicata*, **140**, 181–188.

Li, Y. and Hsing, T. (2010) Uniform convergence rates for nonparametric regression and principal component analysis in functional/longitudinal data. *Annals of Statistics*, **38**, 3321–3351.

McKeague, I. W. and Sen, B. (2010) Fractals with point impact in functional linear regression. *Annals of statistics*, **38**, 2559.

Mentre, F., Mallet, A. and Baccar, D. (1997) Optimal design in random-effects regression models. *Biometrika*, **84**, 429–442.

Müller, H.-G., Carey, J. R., Wu, D., Liedo, P. and Vaupel, J. W. (2001) Reproductive potential predicts longevity of female Mediterranean fruit flies. *Proceedings of the Royal Society B*, **268**, 445–450.

Pearson, J. D., Morrell, C. H., Brant, L. J., Landis, P. K. and Fleg, J. L. (1997) Age-associated changes in blood pressure in a longitudinal study of healthy men and women. *The Journals of Gerontology Series A: Biological Sciences and Medical Sciences*, **52**, M177–M183.

Ramsay, J. and Silverman, B. (2005) *Functional Data Analysis*. Springer, New York.

Rice, J. A. (2004) Functional and longitudinal data analysis: Perspectives on smoothing. *Statistica Sinica*, **14**, 631–648.

Rice, J. A. and Wu, C. O. (2001) Nonparametric mixed effects models for unequally sampled noisy curves. *Biometrics*, **57**, 253–259.

Shock, N. W. *et al.* (1984) Normal human aging: The Baltimore Longitudinal Study of Aging. *NIH Publication*.

Staniswalis, J. G. and Lee, J. J. (1998) Nonparametric regression analysis of longitudinal data. *Journal of the American Statistical Association*, **93**, 1403–1418.

Xiang, D., Qiu, P. and Pu, X. (2013) Nonparametric regression analysis of multivariate longitudinal data. *Statistica Sinica*, **23**, 769–789.

Yao, F., Müller, H.-G. and Wang, J.-L. (2005a) Functional data analysis for sparse longitudinal data. *Journal of the American Statistical Association*, **100**, 577–590.

— (2005b) Functional linear regression analysis for longitudinal data. *Annals of Statistics*, **33**, 2873–2903.

Zagoraiou, M. and Baldi Antognini, A. (2009) Optimal designs for parameter estimation of the Ornstein–Uhlenbeck process. *Applied Stochastic Models in Business and Industry*, **25**, 583–600.

Online Supplement

A.1 Best Linear Predictors

The best linear estimator for $X(t), t \in \mathcal{T}$ in (2) is derived as follows.

Given the observed longitudinal vector $\mathbf{U}$, assume that the best linear estimator for $X(t)$ takes the form

$$B(X(t)|\mathbf{U}) = \mathbf{a} + \mathbf{b^T U}, \tag{28}$$

where a and $\mathbf{b}$ are constant scalar and vector respectively. In this way, $B(X(t)|\mathbf{U})$ minimizes the MSE, defined as

$$\textit{MSE}(g) = \mathrm{E}(X(t) - g(\mathbf{U}))^2 \tag{29}$$

for any estimator $g(\cdot)$. Plugging in the linear presentation, we have

$$\textit{MSE}(B(X(t)|\mathbf{U})) = \mathrm{E}(X(t) - (a + \mathbf{b}^T\mathbf{U}))^{\mathbf{2}}.$$

Taking partial derivatives with respect to a and $\mathbf{b}$ and setting to 0 yields the following two equations for the constants,

$$\frac{\partial}{\partial a}\textit{MSE}(B(X(t)|\mathbf{U}) = 2a + \mathrm{E}(2\mathbf{U}^T\mathbf{b}) - 2\mathrm{E}X(t) = 0, \tag{30}$$

$$\frac{\partial}{\partial \mathbf{b}}\textit{MSE}(B(X(t)|\mathbf{U}) = 2a\mathrm{E}\mathbf{U} + 2\mathrm{E}(\mathbf{U}\mathbf{U}^T\mathbf{b}) - 2\mathrm{E}(X(t)\mathbf{U}) = \mathbf{0}. \tag{31}$$

Observing

$$\begin{aligned}
\mathrm{E}(\mathbf{U}\mathbf{U}^T) &= \mathrm{Cov}(\mathbf{U},\mathbf{U}) + (\mathrm{E}\mathbf{U})(\mathrm{E}\mathbf{U})^T = \mathrm{Cov}(\mathbf{U}) + \boldsymbol{\mu}\boldsymbol{\mu}^T, \\
\mathrm{E}(X(t)\mathbf{U}) &= \mathrm{Cov}(X(t),\mathbf{U}) + \mathrm{E}X(t)\mathrm{E}\mathbf{U} = \mathrm{Cov}(X(t),\mathbf{U}) + \mu(t)\boldsymbol{\mu},
\end{aligned}$$

yields the solutions

$$a = \mu(t) - \mathrm{Cov}(X(t),\mathbf{U})\mathrm{Cov}(\mathbf{U})^{-1}\boldsymbol{\mu}, \quad \mathbf{b} = \mathrm{Cov}(\mathbf{U})^{-1}\mathrm{Cov}(\mathbf{U},X(t)).$$

Plugging in these expressions yields equation (2).

A.2 Local-linear Smoothing for Model Component Estimation

Let $K(\cdot)$ be a symmetric probability density function supported on $[-1,1]$ and $K_h(t) = (1/h)K(t/h)$, where h is a bandwidth or smoothing parameter. Then the mean function, auto-covariance function, cross-covariance surface and measurement error variance σ^2 are estimated as follows, applying the following smoothing steps.

1. The local-linear estimator of the mean function $\mu(t)$ is $\widehat{\mu}_X(t) = \widehat{a}_0$, where

$$(\widehat{a}_0, \widehat{a}_1) = \text{argmin}_{a_0,a_1} \frac{1}{n} \sum_{i=1}^{n} \frac{1}{m_i} \sum_{j=1}^{m_i} \{U_{ij} - a_0 - a_1(t_{ij} - t)\}^2 K_{h_\mu}(t_{ij} - t), \tag{32}$$

with bandwidth h_μ.

2. The local-linear estimator of the cross-covariance function $C(t)$ is $\widehat{C}(t) = \widehat{S}(t) - \widehat{\mu}_X(t)\bar{Y}$, where $\bar{Y} = \frac{1}{n}\sum_{i=1}^{n} Y_i$, $\widehat{S}(t) = \widehat{E}(X(t)Y) = \widehat{a}_0$ and

$$(\widehat{a}_0, \widehat{a}_1) = \text{argmin}_{a_0,a_1} \frac{1}{n} \sum_{i=1}^{n} \frac{1}{m_i} \sum_{j=1}^{m_i} \{U_{ij}Y_i - a_0 - a_1(t_{ij} - t)\}^2 K_{h_S}(t_{ij} - t), \tag{33}$$

with bandwidth h_S.

3. The local-linear estimator of the auto-covariance function $\Gamma(s,t)$ is $\widehat{\Gamma}(s,t) = \widehat{R}(s,t) - \widehat{\mu}_X(s)\widehat{\mu}_X(t)$, where $\widehat{R}(s,t) = \widehat{E}(X(s)X(t)) = \widehat{a}_0$ and

$$(\widehat{a}_0, \widehat{a}_1, \widehat{a}_2) = \text{argmin}_{a_0,a_1,a_2} \frac{1}{n} \sum_{i=1}^{n} [\frac{1}{m_i(m_i-1)} \sum_{k \neq j} \{U_{ij}U_{ik} - a_0 - a_1(t_{ij} - s) - a_2(t_{ik} - t)\}^2 \times K_{h_R}(t_{ij} - s) K_{h_R}(t_{ik} - t)], \tag{34}$$

with bandwidth h_R.

4. Finally, to estimate σ^2, we first obtain an estimate of $V(t) = R(t,t) + \sigma^2$ by a local-linear estimator $\widehat{V}(t) = \widehat{a}_0$, where

$$(\widehat{a}_0, \widehat{a}_1) = \text{argmin}_{a_0,a_1} \frac{1}{n} \sum_{i=1}^{n} \frac{1}{m_i} \sum_{j=1}^{m_i} \{U_{ij}^2 - a_0 - a_1(t_{ij} - t)\}^2 K_{h_V}(t_{ij} - t). \tag{35}$$

Then obtain $\widehat{\sigma}_X^2 = \frac{1}{|\mathcal{T}|} \int_{\mathcal{T}} \{\widehat{V}(t) - \widehat{R}(t,t)\} dt$. To account for boundary effects, one may tailor $\mathcal{T}$ by excluding the boundary area (Yao *et al.* 2005a).

A.3 Maximizing Coefficients of Determination

In this section we show that for trajectory recovery and scalar response regression cases, minimizing either the MISE for trajectory recovery or the MSPE for regression is equivalent to maximizing the proposed coefficients of determination.

Trajectory Recovery. For trajectory recovery, by (5), the MISE as given in (6) becomes

$$\begin{aligned}
\text{MISE}(\mathbf{t}) &= \int_{\mathcal{T}} E\{[X(t) - B(X(t)|\mathbf{U})]^2\}dt \\
&= \int_{\mathcal{T}} E(X(t) - \mu(t) - \boldsymbol{\gamma}(t)^T \boldsymbol{\Gamma}_*^{-1}(\mathbf{U} - \boldsymbol{\mu}))^2 dt \\
&= \int_{\mathcal{T}} \{\text{Var}(X(t)) + \boldsymbol{\gamma}(t)^T \boldsymbol{\Gamma}_*^{-1} \boldsymbol{\Gamma}_* \boldsymbol{\Gamma}_*^{-1} \boldsymbol{\gamma}(t) - 2E[(X(t) - \mu(t))(\mathbf{U} - \boldsymbol{\mu})^T \boldsymbol{\Gamma}_*^{-1} \boldsymbol{\gamma}(t)]\} dt \\
&= \int_{\mathcal{T}} \{\text{Var}(X(t)) - \boldsymbol{\gamma}(t)^T \boldsymbol{\Gamma}_*^{-1} \boldsymbol{\gamma}(t)\} dt. \qquad (36)
\end{aligned}$$

Since the overall variance of the underlying process does not involve the current design points $\mathbf{t}$, minimizing MISE$(\mathbf{t})$ is equivalent to maximizing the term

$$\int_{\mathcal{T}} \boldsymbol{\gamma}(t)^T \boldsymbol{\Gamma}_*^{-1} \boldsymbol{\gamma}(t) dt. \qquad (37)$$

This is exactly the optimization criterion described in (10). Note that the fact that $\int_{\mathcal{T}} \text{Var}(X(t))dt$ is constant implies the equivalence to maximizing R_X^2 in (8).

Scalar Response Regression. Based on (17), the mean squared prediction error (MSPE) given by (18) becomes

$$\begin{aligned}
E[Y - B\{E(Y|X)|\mathbf{U}\}]^2 &= E[Y - \mu_Y - \boldsymbol{\Gamma}_*^{-1}\mathbf{C}^T(\mathbf{U} - \mu)]^2 \\
&= \text{Var}(Y) + \mathbf{C}^T \boldsymbol{\Gamma}_* \boldsymbol{\Gamma}_*^{-1} \boldsymbol{\Gamma}_* \mathbf{C} - 2\mathbf{C}^T \boldsymbol{\Gamma}_*^{-1} \mathbf{C} = \text{Var}(Y) - \mathbf{C}^T \boldsymbol{\Gamma}_*^{-1} \mathbf{C}.
\end{aligned}$$

This implies the equivalence of minimizing MSPE to maximizing the criterion given in (20).

A.4 Asymptotic Results: Assumptions and Proof of Theorem

Following Li and Hsing (2010), let $\eta_{nk} = \left(n^{-1} \sum_{i=1}^n m_i^{-k}\right)^{-1}$, and assume

(A1) For some constants $0 < m < M < \infty$, $m \leq f_T(t) \leq M$ for all $t \in \mathcal{T}$. Further, f_T is differentiable with a bounded derivative.

(A2) The kernel function $K(\cdot)$ is a symmetric probability density function on $[-1, 1]$, and is of bounded variation on $[-1, 1]$. Define $v_2 = \int_{-1}^1 t^2 K(t)dt$.

(A3) $\mu(\cdot)$ is twice differentiable and the second derivative is bounded on $\mathcal{T}$.

(A4) All second-order partial derivatives of $\Gamma(s, t)$ exist and are bounded on $\mathcal{T}^2$.

(A5) $E(|e_{ij}|^{\lambda_\mu}) < \infty$ and $E(\sup_{t\in\mathcal{T}} |X(t)|^{\lambda_\mu}) < \infty$ for some $\lambda_\mu \in (2, \infty)$; $h_\mu \to 0$ and $(h_\mu^2 + h_\mu/\eta_{n1})^{-1}(\log n/n)^{1-2/\lambda_\mu} \to 0$ as $n \to \infty$.

(A6) $E(|e_{ij}|^{2\lambda_R}) < \infty$ and $E(\sup_{t\in\mathcal{T}}|X(t)|^{2\lambda_R}) < \infty$ for some $\lambda_R \in (2,\infty)$; $h_R \to 0$ and $(h_R^4 + h_R^3/\eta_{n1} + h_R^2/\eta_{n2})^{-1}(\log n/n)^{1-2/\lambda_R} \to 0$ as $n \to \infty$.

(A7) $E(|e_{ij}|^{2\lambda_V}) < \infty$ and $E(\sup_{t\in\mathcal{T}}|X(t)|^{2\lambda_V}) < \infty$ for some $\lambda_V \in (2,\infty)$; $h_V \to 0$ and $(h_V^2 + h_V/\eta_{n1})^{-1}(\log n/n)^{1-2\lambda_V} \to 0$ as $n \to \infty$.

Proof of the Theorem. Based on the assumptions (A1)-(A7) and the smoothing methods we use (see (32), (34), (33)) for the estimates of $\mu(t), \Gamma(s,t), C(t)$, uniform convergence and rates of convergence are shown in (Li and Hsing 2010). If $h_R^2 \lesssim h_\mu \lesssim h_R$, $h_V + (\log n/n)^{1/3} \lesssim h_R \lesssim h_V^2 n/\log n$, we have, as $n \to \infty$,

$$\sup_{t\in\mathcal{T}}|\widehat{\mu}_X(t) - \mu(t)| = O(h_\mu^2 + \{\log n/(nh_\mu)\}^{1/2}) \ \ a.s.,$$
$$\sup_{t\in\mathcal{T}}|\widehat{C}(t) - C(t)| = O(h_S^2 + \{\log n/(nh_S)\}^{1/2}) \ \ a.s.,$$
$$\sup_{s,t\in\mathcal{T}}|\widehat{\Gamma}(s,t) - \Gamma(s,t)| = O(h_R^2 + \{\log n/(nh_R^2)\}^{1/2}) \ \ a.s.$$
$$\widehat{\sigma}_X^2 - \sigma^2 = O(h_R^2 + \{\log n/(nh_R)\}^{1/2}) \ \ a.s.$$

Set $\theta_n = h_R^2 + \{\log n/(nh_R^2)\}^{1/2}$. It is easy to see that $h_R^2 + \{\log n/(nh_R)\}^{1/2} = O(\theta_n)$. Since p is fixed and finite, $\mathcal{T}$ is compact, and that X is $\mathcal{L}^2$ process, we have that $\Gamma(s,t)$ is finite for all $s,t \in \mathcal{T}$. As one assumption is $h_R^2 \lesssim h_\mu \lesssim h_R$, for some $\delta_0 > 0$

$$\sup_{t\in\mathcal{T},\mathbf{t}\in\mathcal{T}(\delta_0)}|\widehat{\boldsymbol{\gamma}}(t)^T\widehat{\boldsymbol{\Gamma}}_*^{-1}\widehat{\boldsymbol{\gamma}}(t) - \boldsymbol{\gamma}(t)^T\boldsymbol{\Gamma}_*^{-1}\boldsymbol{\gamma}(t)| = O(\theta_n) \ \ a.s.$$

as $n \to \infty$. So for all $\mathbf{t} \in \mathcal{T}(\delta_0)$, writing $g(\mathbf{t}) = \int_{\mathcal{T}} \boldsymbol{\gamma}(t)^T\boldsymbol{\Gamma}_*^{-1}\boldsymbol{\gamma}(t)dt$, and $\widehat{g}(\mathbf{t}) = \int_{\mathcal{T}} \widehat{\boldsymbol{\gamma}}(t)^T\widehat{\boldsymbol{\Gamma}}_*^{-1}\widehat{\boldsymbol{\gamma}}(t)dt$, we readily have

$$|\widehat{g}(\mathbf{t}) - g(\mathbf{t})| \le \int_{\mathcal{T}}|\widehat{\boldsymbol{\gamma}}(t)^T\widehat{\boldsymbol{\Gamma}}_*^{-1}\widehat{\boldsymbol{\gamma}}(t) - \boldsymbol{\gamma}(t)^T\boldsymbol{\Gamma}_*^{-1}\boldsymbol{\gamma}(t)|dt = O(\theta_n) \ \ a.s.,$$

and therefore,

$$\sup_{\mathbf{t}\in\mathcal{T}(\delta_0)}|\widehat{g}(\mathbf{t}) - g(\mathbf{t})| \xrightarrow{n\to\infty} 0 \ \ a.s.$$

For trajectory recovery, the true optimal design is $\mathbf{t}_X^* = \arg\max_{\mathbf{t}\in\mathcal{T}(\delta_0)} g(\mathbf{t})$, and its estimate is $\widehat{\mathbf{t}}_X^* = \arg\max_{\mathbf{t}\in\mathcal{T}(\delta_0)} \widehat{g}(\mathbf{t})$. We claim that $|g(\mathbf{t}_X^*) - \widehat{g}(\widehat{\mathbf{t}}_X^*)| = O(\theta_n)$ a.s., as when

$g(\mathbf{t}_X^*) \ge \widehat{g}(\widehat{\mathbf{t}}_X^*)$, $0 \le g(\mathbf{t}_X^*) - \widehat{g}(\widehat{\mathbf{t}}_X^*) \le g(\mathbf{t}_X^*) - \widehat{g}(\mathbf{t}_X^*) = O(\theta_n) \ \ a.s.$, and

$g(\mathbf{t}_X^*) \le \widehat{g}(\widehat{\mathbf{t}}_X^*)$, $0 \le \widehat{g}(\widehat{\mathbf{t}}_X^*) - g(\mathbf{t}_X^*) \le \widehat{g}(\widehat{\mathbf{t}}_X^*) - g(\widehat{\mathbf{t}}_X^*) = O(\theta_n) \ \ a.s.$

By smoothness of $\Gamma(s,t)$, we have that $g : \mathcal{T}(\delta_0) \to \mathbb{R}$ is a continuous function. Now,

$$|g(\mathbf{t}_X^*) - g(\widehat{\mathbf{t}}_X^*)| \le |g(\mathbf{t}_X^*) - \widehat{g}(\widehat{\mathbf{t}}_X^*)| + |\widehat{g}(\widehat{\mathbf{t}}_X^*) - g(\widehat{\mathbf{t}}_X^*)| = O(\theta_n) \ \ a.s.$$

which implies that

$$\widehat{\mathbf{t}}_X^* \xrightarrow{n\to\infty} \mathbf{t}_X^* \ \ a.s.$$

Moreover, with a mean value $\boldsymbol{\xi}$,

$$|g(\mathbf{t}_X^*) - g(\widehat{\mathbf{t}}_X^*)| = \frac{1}{2}(\widehat{\mathbf{t}}_X^* - \mathbf{t}_X^*)^T g^{(2)}(\boldsymbol{\xi})(\widehat{\mathbf{t}}_X^* - \mathbf{t}_X^*).$$

Using the assumption that g is locally concave around $\mathbf{t}_X^*$ and the convergence derived above, we have

$$||\widehat{\mathbf{t}}_X^* - \mathbf{t}_X^*||_p = O(\theta_n^{1/2}) \ \ a.s.$$

where $||\mathbf{t} - \mathbf{s}||_p = \max_{1\le j\le p} |\mathbf{t}_{(j)} - \mathbf{s}_{(j)}|$, where $\mathbf{t}_{(j)}$, $\mathbf{s}_{(j)}$ are the j-th order statistics of designs $\mathbf{t}$ and $\mathbf{s}$.

The proof for optimal regression designs is analogous.

A.5 Comparison between Optimal Designs and Random Designs

We compare selected optimal designs with "random designs" where measurement locations are uniformly distributed over the domain, because in the sparse case there are no special designs that are representative for the data and equidistant designs are not commonly (in practice almost never) observed. For trajectory recovery of longitudinal data, it can be expected that the optimal designs are somewhat evenly spaced if the data show even distribution of variation across subjects, for example if random curvature or trends are modest.

This is the case for the BLSA data and the bike sharing data. For scalar response regression, the situation is more complex, as the covariance function matters, and again there are no a priori special designs in the sparse random design case. Therefore it is best to base comparisons on the entire spectrum of designs that may be encountered, which is well approximated with the space of random designs. Therefore the best comparison is indeed with "random designs" to get a general idea of how the selected optimal designs work.

A.6 Additional Example: Bike Sharing

Bike sharing systems increasingly replace traditional bike rentals. In such systems the entire process of rental and return is automatized and data about the rentals are automatically generated. The data we use are described in Fanaee-T and Gama (2013) and include records of the hourly number of

bike rentals for 2011 and 2012 from the Capital Bikeshare system, Washington D.C. We consider the trajectories corresponding to hourly numbers of bike rentals for all Saturdays and aim at finding the optimal design points to recover these bike rental trajectories.

The bike rental data are densely and regularly measured, with 24 measurements per day, of which we randomly chose 3 to 5 measurements to create sparse designs, thus rendering the pilot data sparse. The ridge parameter needed for implementation of the method is obtained by modified cross-validation. We focus on Saturday bike rentals and compare the recovered trajectories based on the selected optimal designs that are constructed from the sparsified pilot data with the actually observed curves. To assess the performance of the proposed optimal designs in comparison to random designs, we include data for 80 days in the training and for 22 days in the testing sample. Figure 6 shows the Spaghetti plot for the training sample. The goal is to choose optimal design points for designs with $p = 3$ measurements.

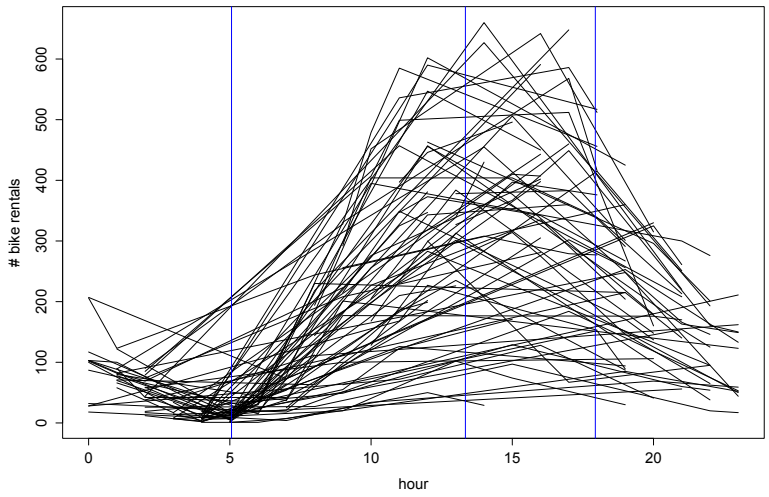

Figure 6: Spaghetti plot for training sample of bike sharing data. The blue vertical bars denote the locations of the optimal design points for $p = 3$ for recovering daily bike rental trajectories.

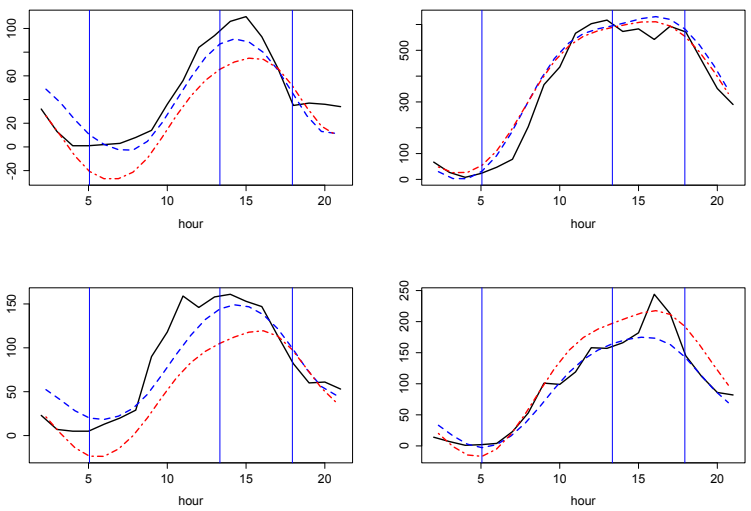

Figure 7: Trajectory recovery results for four randomly chosen days from the test sample for the bike sharing data with sparse pilot measurements for $p = 3$. Black curves stand for the underlying dense data. Blue dashed and red dotted curves are trajectory estimates using optimal designs and random designs, respectively. The vertical blue lines are the locations for the optimal design points.

Randomly selected subjects from the test set and their recovered trajectories are illustrated in Figure 7. Relative ARE (26) for the test sample using optimal designs is 0.264, compared to 0.321 for random designs of median performance for $p = 3$. In Figure 7, the optimal designs are also seen to perform better than median performance random designs. The selected optimal design points are located at 5am, 1:30pm and 6pm during the day. These three points are located around time points corresponding to morning valley, noon peak and afternoon decrease of the bike renting profile for each day. These three points thus likely reflect key features of the daily bike renting profiles and therefore it is not surprising that their locations constitute the optimal design.

A.7 Additional Discussion on Simulation Studies

In our simulation studies, we also compared sequential design selection and exhaustive search and the effect of the ridge parameter σ^2_{new} on the performance of optimal designs. The results are summarized in Table 3 and Figure 8.

Table 3: Investigating Sequential Optimization on Optimal Designs in Simulations, in Terms of Mean ARE and Relative ARE (26) (in brackets) for Trajectory Recovery and APE and Relative APE (27) (in brackets) for Response Prediction in a Functional Linear Model based on 100 simulations.

	Mean ARE (ARE*)		Mean APE (APE*)	
Sequential	Dense	Sparse	Dense	Sparse
$p = 2$	$1.74(.092)$	$1.83(.101)$	$2.91(.234)$	$9.23(.745)$
$p = 3$	$1.37(.072)$	$1.59(.088)$	$2.4(.193)$	$7.85(.630)$
$p = 4$	$1.02(.054)$	$1.34(.074)$	$2.18(.175)$	$7.48(.601)$
$p = 5$	$1.01(.053)$	$1.32(.073)$	$1.91(.154)$	$6.71(.542)$
$p = 6$	$0.91(.048)$	$1.24(.069)$	$1.68(.136)$	$6.29(.508)$

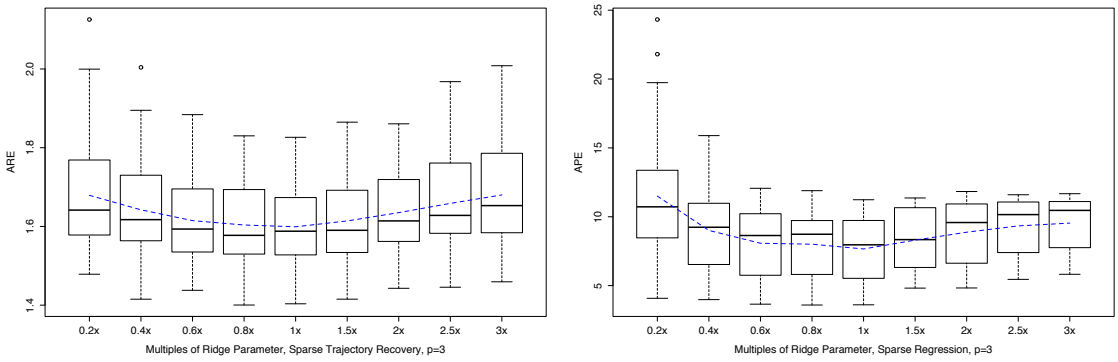

Figure 8: Simulation results for the effect of the ridge parameter σ^2_{new} for 100 simulation runs. The left (right) panel depicts boxplots of ARE (APE) from sparsely observed simulation data by using multiples of the optimal ridge parameter selected by cross validation for trajectory recovery (scalar response regression) for $p = 3$. Blue dotted lines denote the mean error across the different ridge parameters.

Table 3 summarizes the simulation results of optimal designs from sequential selection algorithm for both dense and sparse scenarios with $p = 2, 3, 4, 5, 6$ based on 100 simulations. The performance of sequential optimization is comparable to that of exhaustive search in all simulation scenarios, indicating that in the simulation scenario only a minor loss in optimality occurs in exchange for a substantial reduction in computing time.

Ridge parameter also affects the performance of optimal designs. Figure 7 shows ARE (APE) based on different multiples of selected ridge parameter via cross validation. For both simulation scenarios, the error is not very sensitive to the choice of ridge parameter, and similar results are observed for different p as well as for densely observed data. In practice, the candidate set from which ridge parameters are selected can be chosen heuristically, where the selected outcome should not be located at the boundary of the candidate set. The selection is relatively fast when combined with the sequential search algorithm for optimal designs.

A.8 Additional Discussion

For some applications where sequential design selection is desirable, given that our proposed method relies on estimates of mean and covariance of the functional data, we could continuously update these estimates when additional measurements for one or more subjects become available. This would be done by combining the new measurements with the current pilot data and then recalculating the optimal design. The resulting optimal designs would then be applied for the next batch of subjects, and one could iterate this scheme. Though numerically possible, we do not recommend updating the optimal design given a subset of measurements for a particular subject, because it is less practical to reschedule subject visits in longitudinal studies as would be advisable when the design has changed, and also mean and covariance surface are unlikely to change much when adding the measurements of a single subject or a subset of these measurements.

Another extension of our proposed method is to select optimal designs reaching a specified predictive accuracy threshold at a location in the domain which is as early as possible, so that predictions of scalar outcomes or the remaining random trajectory can then be computed relatively early. This would allow a longitudinal study to run for a shorter time, which can lead to faster results and cost savings.

The predictive accuracy threshold in this setting can be chosen as the coefficient of determination R^2 (subscripts X and Y suppressed for convenience), which is scale-invariant and applicable for both functional and longitudinal pilot studies. Suppose that the functional support is $\mathcal{T} = [0, T]$, and define $R^2(t)$ for $t \in (0, T]$ as the coefficient of determination attained by the optimal design, selected with the proposed method, but using data only on the time domain $[0, t]$. The function $R^2(t)$ increases with t because the candidate set of designs becomes larger as t increases. With a user-specified threshold parameter $\alpha \in (0, 1)$, the final outcome is the optimal design that corresponds to the optimal design

obtained from data on the interval $[0, t_\alpha]$, where

$$t_\alpha = \min\{t \in (0, T] : R^2(t) \geq (1 - \alpha)R^2(T)\}.$$

This version of constructing optimal designs can be appealing for certain applications, since it could potentially save a lot of resources, with an acceptable decrease in predictive performance.

A.9 Additional Figures

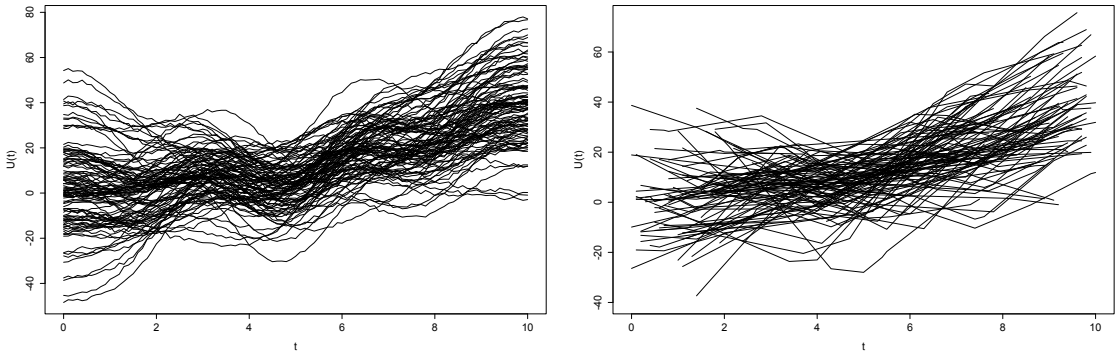

Figure 9: Training samples for one simulation run. Left panel depicts a sample of size $n = 100$ densely sampled functional data (scenario 1) and right panel is the spaghetti plot of a sample of 100 simulated training data based on a sparse functional sampling design (scenario 2).

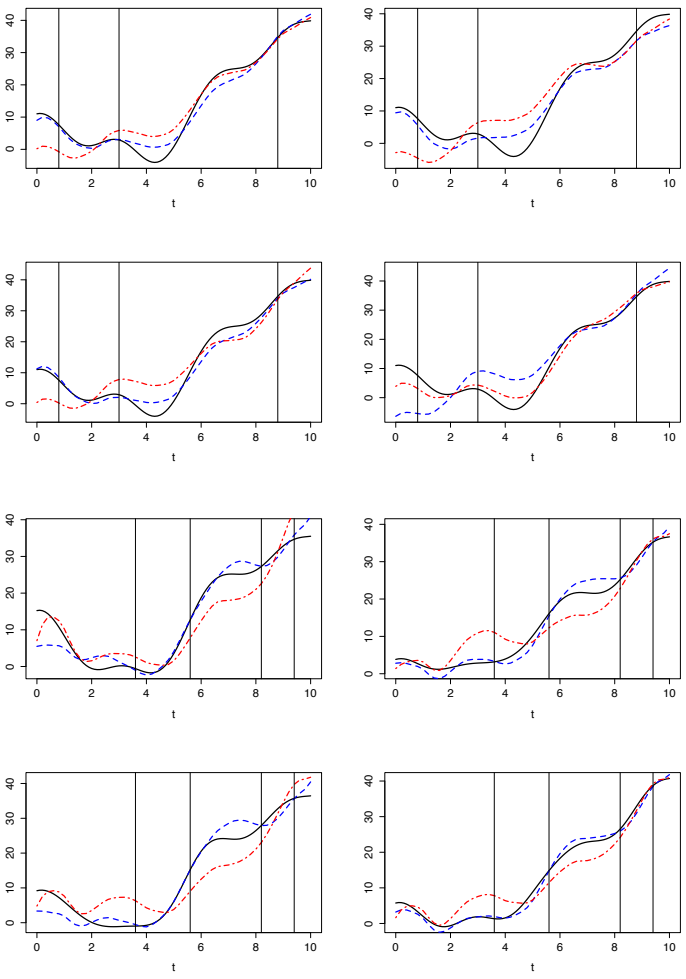

Figure 10: Simulation results for sparsely sampled pilot data: Recovered trajectories for eight randomly chosen subjects from the test sample, based on 3 (upper four panels) or 4 (lower four panels) design points. Included are in each panel the true underlying curves (black solid curves), estimated trajectories from random designs with 3 (4) support points (red dash-dotted curves), for the random design with median performance, the estimated trajectories for the optimal designs (blue dashed curves), and the optimal design point locations (black vertical lines).

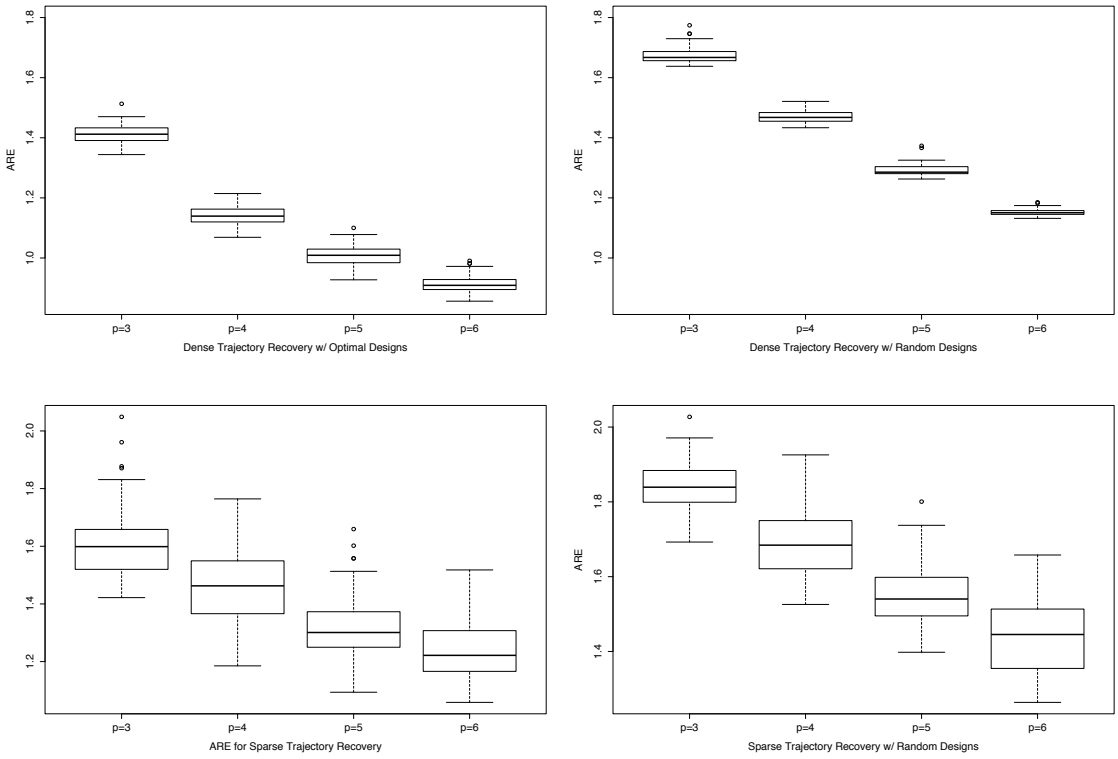

Figure 11: Boxplots of ARE from 100 simulation runs for trajectory recovery for both dense and sparse scenarios with sequential search, for $p = 3, 4, 5, 6$. The first (second) row corresponds to a dense (sparse) scenario; left (right) column is for the optimal design (median performance random design).